

LA-UR-13-28913

Approved for public release; distribution is unlimited.

Title: A Hybrid Deterministic/Monte Carlo Method for Solving the k-Eigenvalue Problem with a Comparison to Analog Monte Carlo Solutions

Author(s): Willert, Jeffrey A.
Knoll, Dana A.
Park, Hyeongkae
Kelley, Carl T.

Intended for: Transport Theory and Statistical Physics

Issued: 2013-11-19



Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

A Hybrid Deterministic/Monte Carlo Method for Solving the k -Eigenvalue Problem with a Comparison to Analog Monte Carlo Solutions

Jeffrey Willert^a, C.T. Kelley^{1b}, D. A. Knoll^a, H. Park^a

^a*Theoretical Division, MS B216, Los Alamos National Laboratory, Los Alamos, NM 87545*

^b*Department of Mathematics, North Carolina State University, Raleigh, NC 27695-8205*

Abstract

In this paper we present a hybrid deterministic/Monte Carlo algorithm for computing the dominant eigenvalue/eigenvector pair for the neutron transport k -eigenvalue problem in multiple space dimensions. We begin by deriving the Nonlinear Diffusion Acceleration method [2, 5] for the k -eigenvalue problem. We demonstrate that we can adapt the algorithm to utilize a Monte Carlo simulation in place of a deterministic transport sweep. We then show that the new hybrid method can be used to solve a two-group, two-dimensional eigenvalue problem. The hybrid method is competitive with analog Monte Carlo in terms of the number of particle flights required to compute the eigenvalue, however it produces a much less noisy eigenvector and fission source distribution. Furthermore, we show that we can reduce error induced by the discretization of the low-order system by an appropriate refinement of the mesh.

Keywords: Neutron Transport, k -Eigenvalue Problem, Hybrid Methods, Moment-Based Acceleration

¹The work of this author has been partially supported by the Consortium for Advanced Simulation of Light Water Reactors (www.casl.gov), an Energy Innovation Hub (<http://www.energy.gov/hubs>) for Modeling and Simulation of Nuclear Reactors under U.S. Department of Energy Contract No. DE-AC05-00OR22725.

1. Introduction

In recent years, the use and understanding of Moment-Based Acceleration techniques (MBAs) has increased dramatically for the solution the transport problems [2, 3, 4, 5, 11]. These accelerations use a consistent system of moment equations to accelerate the solution to the original transport problem. These methods have been extended for the solution of the neutron transport k -eigenvalue problem and have been shown to be extremely efficient. On the other hand, an active area of research in the neutron transport community is the use of hybrid deterministic/Monte Carlo methods for the solution of the neutron transport equation [9, 8, 7, 3, 4, 11]. The idea with these hybrid methods is to remove all sources of discretization error while simultaneously accelerating the Monte Carlo simulation. While the authors of [7] seek to only accelerate the convergence of the fission source, our algorithm simultaneously attempts to accelerate both the scattering and fission sources. Purely deterministic methods suffer from spatial discretization error, ray effects and errors in reaction rates from incorporating the multi-group approximation. In addition, Monte Carlo methods are generally highly parallelizable which makes them strong candidates for efficient implementation on current and emerging computer architectures.

In this paper, we demonstrate that we can modify a Nonlinear Diffusion Acceleration (NDA) based method to use a Monte Carlo transport sweep. We derive the method and discuss the theoretical benefits of using a Monte Carlo transport sweep to solve the k -eigenvalue problem in multiple dimensions [11]. At this point, we discuss relevant details of the implementation and demonstrate the new method on 2-group, 2-D eigenvalue problem. To conclude we provide a head-to-head comparison of the hybrid method against an analog Monte Carlo implementation, before discussing the advantages and disadvantages of using the hybrid algorithm.

2. Nonlinear Diffusion Acceleration for the Transport Criticality Problem

We are interested in solving the multi-group formulation of the neutron transport eigenvalue problem with isotropic scattering,

$$\hat{\Omega} \cdot \nabla \psi_g(\hat{\Omega}, \vec{r}) + \Sigma_{t,g} \psi_g(\hat{\Omega}, \vec{r}) = \frac{1}{4\pi} \left[\sum_{g'=1}^G \Sigma_s^{g' \rightarrow g} \phi_{g'}(\vec{r}) + \frac{\chi_g}{k_{eff}} \sum_{g'=1}^G \nu \Sigma_{f,g'} \phi_{g'}(\vec{r}) \right], \quad (1)$$

in which ψ_g is the group angular flux and $\phi_g = \int_{4\pi} \psi_g d\hat{\Omega}$ is the group scalar flux for groups $g = 1, \dots, G$. $\Sigma_{t,g}$, $\Sigma_s^{g' \rightarrow g}$, and $\Sigma_{f,g}$ are the total, in scattering and fission cross-sections for group g , respectively. Furthermore, χ_g denotes the multi-group fission spectrum, ν is the mean number of neutrons emitted per fission event and k_{eff} is the dominant eigenvalue.

In order to simplify future discussion, we'll represent the eigenvalue problem (Eq. 1) in operator notation,

$$\mathcal{L}\Psi = \frac{1}{4\pi} \left[\mathcal{S} + \frac{1}{k_{eff}} \mathcal{F} \right] \Phi,$$

in which

$$\begin{aligned} \mathcal{L} &= \hat{\Omega} \cdot \nabla + \Sigma_t \\ \mathcal{S} &= \Sigma_s \\ \mathcal{F} &= \chi \nu \Sigma_f. \end{aligned}$$

In this setting, Ψ and Φ represent vectors of group angular and scalar fluxes, respectively,

$$\begin{aligned} \Psi &= [\psi_1, \psi_2, \dots, \psi_G], \\ \Phi &= [\phi_1, \phi_2, \dots, \phi_G]. \end{aligned}$$

2.1. Deriving the NDA Low-Order Equation

As in [2, 3, 4, 5, 11, 12], we begin by computing the zeroth angular moment of Equation 1,

$$\nabla \cdot \vec{J}_g(\vec{r}) + (\Sigma_{t,g} - \Sigma_s^{g \rightarrow g}) \phi_g(\vec{r}) = \sum_{g' \neq g} \Sigma_s^{g' \rightarrow g} \phi_{g'}(\vec{r}) + \frac{\chi_g}{k_{eff}} \sum_{g'=1}^G \nu \Sigma_{f,g'} \phi_{g'}(\vec{r}), \quad (2)$$

in which the current, \vec{J} is defined as the first angular moment,

$$\vec{J}_g = \int_{4\pi} \hat{\Omega} \psi_g(\hat{\Omega}, \vec{r}) d\hat{\Omega}. \quad (3)$$

We refer to Equation 2 as the neutron balance equation. As in [2, 3, 4, 5, 11, 12], we use a Fick's law plus drift term closure relationship for the current,

$$\vec{J}_g = -\frac{1}{3\Sigma_{t,g}} \nabla \phi_g + \hat{D}_g \phi_g. \quad (4)$$

Substituting Equation 4 into Equation 2 yields the NDA low-order system,

$$\begin{aligned} \nabla \cdot \left[-\frac{1}{3\Sigma_{t,g}} \nabla \phi_g + \hat{D}_g \phi_g \right] + (\Sigma_{t,g} - \Sigma_s^{g \rightarrow g}) \phi_g(\vec{r}) = \\ \sum_{g' \neq g} \Sigma_s^{g' \rightarrow g} \phi_{g'}(\vec{r}) + \frac{\chi_g}{k_{eff}} \sum_{g'=1}^G \nu \Sigma_{f,g'} \phi_{g'}(\vec{r}). \end{aligned} \quad (5)$$

The consistency term \hat{D}_g is computed using high-order quantities of the scalar flux, ϕ_g^{HO} , and current, \vec{J}_g^{HO} , via

$$\hat{D}_g = \frac{\vec{J}_g^{HO} + \frac{1}{3\Sigma_{t,g}} \nabla \phi_g^{HO}}{\phi_g^{HO}}. \quad (6)$$

These high-order moments are computed directly via integration of the high-order angular flux, which is the results of a single high-order transport sweep. That is,

$$\begin{aligned} \Psi^{HO} &= \frac{1}{4\pi} \mathcal{L}^{-1} \left[\mathcal{S} + \frac{1}{k_{eff}} \mathcal{F} \right] \Phi, \\ \Phi^{HO} &= \int \Psi^{HO} d\hat{\Omega}, \\ \vec{J}^{HO} &= \int \hat{\Omega} \Psi^{HO} d\hat{\Omega}. \end{aligned}$$

We will express Equation 5 in operator notation as well,

$$\mathcal{D}\Phi = (\mathcal{S}_U + \mathcal{S}_L) \Phi + \frac{1}{k_{eff}} \mathcal{F}\Phi, \quad (7)$$

in which

$$\begin{aligned} \mathcal{D}_g \Phi &= \nabla \cdot \left[-\frac{1}{3\Sigma_{t,g}} \nabla + \hat{D}_g \right] \phi_g + (\Sigma_{t,g} - \Sigma_s^{g \rightarrow g}) \phi_g, \\ (\mathcal{S}_{U,g} + \mathcal{S}_{L,g}) \Phi &= \sum_{g' \neq g} \Sigma_s^{g' \rightarrow g} \phi_{g'}. \end{aligned}$$

2.2. NDA-Based Eigenvalue Methods

When we apply NDA to the neutron transport k -eigenvalue problem [2, 5] it yields a set of algorithms in which we alternate between executing high-order transport sweeps and solving the low-order eigenvalue problem (Equation 5). These methods differ only in the way in which the low-order eigenvalue problem is solved. We have described the methods formally in Algorithm 1

Algorithm 1 NDA

Nonlinear Diffusion Acceleration

Compute initial iterate $\Phi^{(0)}$, initial eigenvalue approximation k^0 . Set iteration counter $m = 0$.

while $|k^m - k^{m-1}| > \tau$ **do**

 Update counter, $m = m + 1$.

 Execute transport sweep and compute new consistency term

$$\Psi^{(m)} = \frac{1}{4\pi} \mathcal{L}^{-1} \left(\mathcal{S} + \frac{1}{k^{m-1}} \mathcal{F} \right) \Phi^{(m-1)}, \quad (8)$$

$$\Phi^{HO} = \int \Psi^{(m)} d\hat{\Omega}, \quad (9)$$

$$\vec{J}^{HO} = \int \hat{\Omega} \Psi^{(m)} d\hat{\Omega}, \quad (10)$$

$$\hat{D}^{(m)} = \frac{\vec{J}^{HO} + \frac{1}{3\Sigma_t} \nabla \Phi^{HO}}{\Phi^{HO}}. \quad (11)$$

Solve the low-order eigenvalue problem for $\Phi^{(m)}$ and k^m

$$(\mathcal{D}^{(m)} - \mathcal{S}_U - \mathcal{S}_L) \Phi^{(m)} = \frac{1}{k^{(m)}} \mathcal{F} \Phi^{(m)}. \quad (12)$$

end while

We generally choose to solve the low-order eigenvalue problem using Newton's method, which is known as the NDA-NCA method for computing the dominant eigenvalue [5, 3, 11, 12]. For the purpose of this paper, we will not concern ourselves with details regarding the solution to the low-order eigenvalue problem. For a more in-depth treatment of the solution to the low-order eigenvalue problem, see [12].

3. Hybrid NDA-NCA

Up until this point, we have not specified a spatial or angular discretization for the high-order problem. This was done intentionally to demonstrate that it is unnecessary to use a deterministic transport sweep. When these algorithms were originally formulated, S_n quadratures were used alongside a step characteristics or linear-discontinuous Galerkin spatial discretization. However, in this paper, we utilize a “Monte Carlo transport sweep” to recover Φ^{HO} and \vec{J}^{HO} . This sweep is described in Algorithm 2.

Algorithm 2 Monte Carlo Transport Sweep

Monte Carlo Transport Sweep

Input current approximation to the scalar flux, $\Phi^{(m-1)}$.

Build a scattering and fission-free fixed source problem:

$$\mathcal{L}\Psi = \mathcal{Q} \tag{13}$$

in which

$$\mathcal{Q} = \frac{1}{4\pi} \left(\mathcal{S} + \frac{1}{k^{(m-1)}} \mathcal{F} \right) \Phi^{(m-1)}$$

Solve Equation 13 via Monte Carlo simulation to recover Φ^{HO} and \vec{J}^{HO} .

It is important to note the difference between the Monte Carlo transport sweep described in Algorithm 2 and a standard Monte Carlo simulation. In a standard Monte Carlo simulation, we must simulate the entire life of a neutron, including scattering, absorption, streaming, fission, etc. However, in the Monte Carlo transport sweep, the simulation is used to solve a scattering- and fission-free problem. Each particle is born at a specific location, streams some distance and is immediately absorbed at the first interaction location. The fact that each neutron undergoes a single particle flight makes the Monte Carlo transport sweep far less expensive to execute and minimizes the logic involved in simulating the particle history [3, 4].

4. Implementation

In this section we discuss the relevant details of the implementation of the hybrid NDA-NCA algorithm (NDA-NCA(MC)). For each problem, we

choose our initial iterate by solving the diffusion eigenvalue problem for Φ and k . For the problems we are interested in, this provides a much better starting point than using a flat source.

It should be readily apparent to the reader that the accuracy of the simulation is related to the number of particles used in each Monte Carlo transport sweep. It is well known, see [17], that the standard deviation in a Monte Carlo simulation decreases proportional to $\frac{1}{\sqrt{N}}$ in which N is the number of particle histories. Therefore, we can increase the number of particles per transport sweep until we have achieved the desired level of accuracy (or until the computational costs become too burdensome). However, at early stages of the iteration, the level of accuracy necessary to advance the iteration may be significantly lower than at the terminal stage. Willert et. al. showed that for fixed-source computations it is most efficient to begin with a relatively small number of particles for the first iteration and increase the value by some factor at each transport sweep [13]. Therefore, for these two-dimensional test problems, we begin by using 10^9 particles per transport sweep and increase this number by a factor of 2 at each iteration.

Also, as was previously mentioned, we utilize Newton's method to solve the low-order eigenvalue problem. We do this by nonlinearly eliminating the eigenvalue,

$$F(\Phi) = (\mathcal{D}^{(m)} - \mathcal{S}_U - \mathcal{S}_L) \Phi - \frac{1}{k(\Phi)} \mathcal{F} \Phi. \quad (14)$$

in which

$$k(\Phi) = \int \mathcal{F} \Phi dV.$$

This equation used for the elimination of k is justified using a proper normalization,

$$1 = \frac{\int \mathcal{F} \Phi^{(m-1)} dV}{k^{(m-1)}}.$$

We precondition the Newton-Krylov iteration using the inverse of the operator \mathcal{M} formed by evaluating k at the current approximation of the eigenvalue,

$$\mathcal{M} = (\mathcal{D}^{(m)} - \mathcal{S}_U - \mathcal{S}_L) - \frac{1}{k(\Phi^{HO})} \mathcal{F}.$$

To ensure convergence, we employ a standard Newton’s method line-search [1] as well as a “physics-based” line-search. We use our knowledge of the positivity of the dominant eigenvector to demand that Φ remain strictly positive everywhere in the domain. The combination of these two line-searches allows us to converge to the dominant eigenpair without difficulty. While we have no definitive theory which demands that this iteration converge to the dominant eigenpair, we have yet to encounter a case in which it does not.

It is also very important to note that this algorithm is highly scalable. In Figure 1 we demonstrate the strong scaling of the algorithm using both standard track length tallies and the more-efficient, balance-preserving continuous energy deposition (CED) tallies [3, 16]. Using standard track length tallies we see a weak scaling efficiency of roughly 95% and using the CED Tallies we demonstrate a weak scaling efficiency of nearly 99%. Furthermore, Willert et. al. demonstrated the efficiency with which these Monte Carlo transport sweeps could be implemented on general-purpose, graphical processing units (GP-GPUs) for a similar problem [19].

The low-order problem is discretized using a finite-difference discretization. At each iteration, the low-order problem provides a constant-cell source from each cell to the Monte Carlo transport sweep. We acknowledge that this may be a deficiency of the algorithm and will consider this as an area of future work. By using a higher-order discretization or applying shape functions within each cell, we may reduce the effects of the low-order discretization. For now, these discretization errors can be alleviated by refining the low-order mesh. It is important to note, however, refining the mesh does not effect the overall cost of the algorithm in a dramatic way. The cost of the low-order solve increases slightly, but this increase is inconsequential compared to the cost of the high-order Monte Carlo transport sweep.

5. Numerical Results

In Figure 2 we display the material layout for the LRA-BWR problem [14, 15]. This is a five-material problem in which materials 1 through 4 consist of fissionable material and material 5 is a reflector region.

As previously stated, the initial Monte Carlo transport sweep utilizes 10^9 particle histories and each subsequent sweep uses a factor of 2 times the previous. We must recall that the low-order spatial mesh may bias the solution due to the discretization error. For this reason, we will solve the

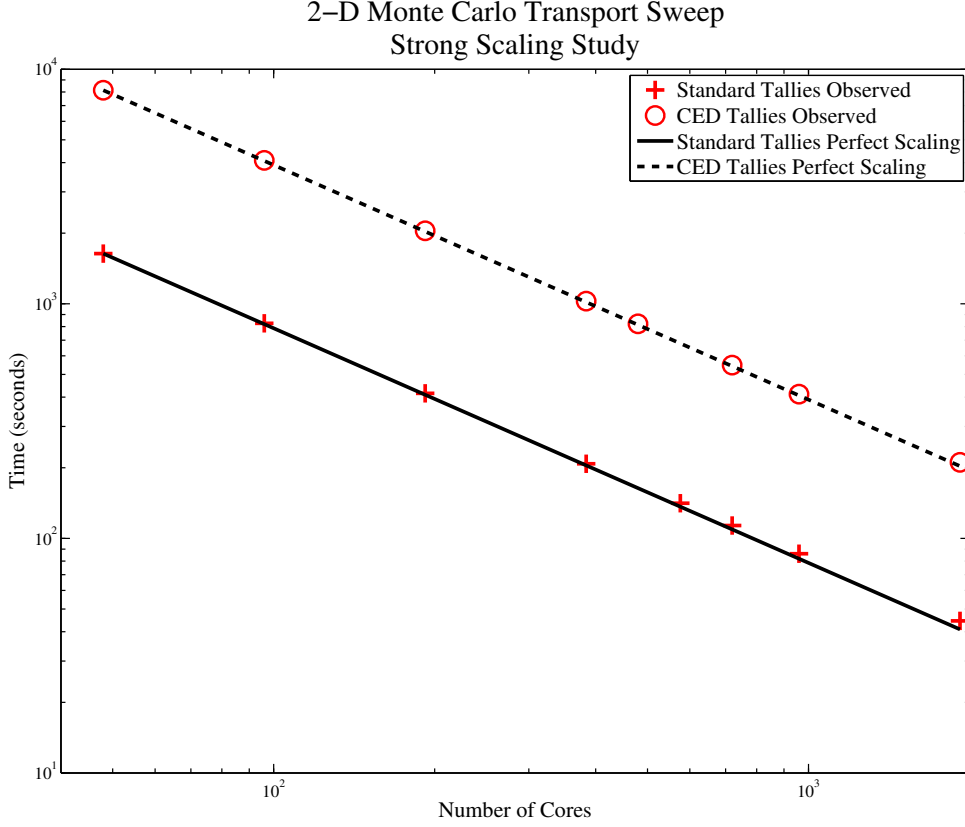


Figure 1: The NDA-NCA(MC) algorithm is highly scalable.

problem on a series of meshes in order to demonstrate that we should be able to compute the true eigenvalue using the appropriate level of mesh refinement.

For this problem, we have computed a reference eigenvalue $k_{ref} = 1.00144$ using an analog Monte Carlo simulation. On a 165×165 mesh using an S_{16} quadrature, the deterministic method computes $k_{det} = .99896$ using a step-characteristics discretization. On the same low-order spatial mesh, the hybrid method computes $k_{hyb} = 1.00022$. Clearly, both the deterministic method and hybrid method have some level of discretization error. With the deterministic method, this error comes in the form of both spatial truncation error and angular discretization error. With the hybrid method this error comes only from the discretization error of the low-order problem. The

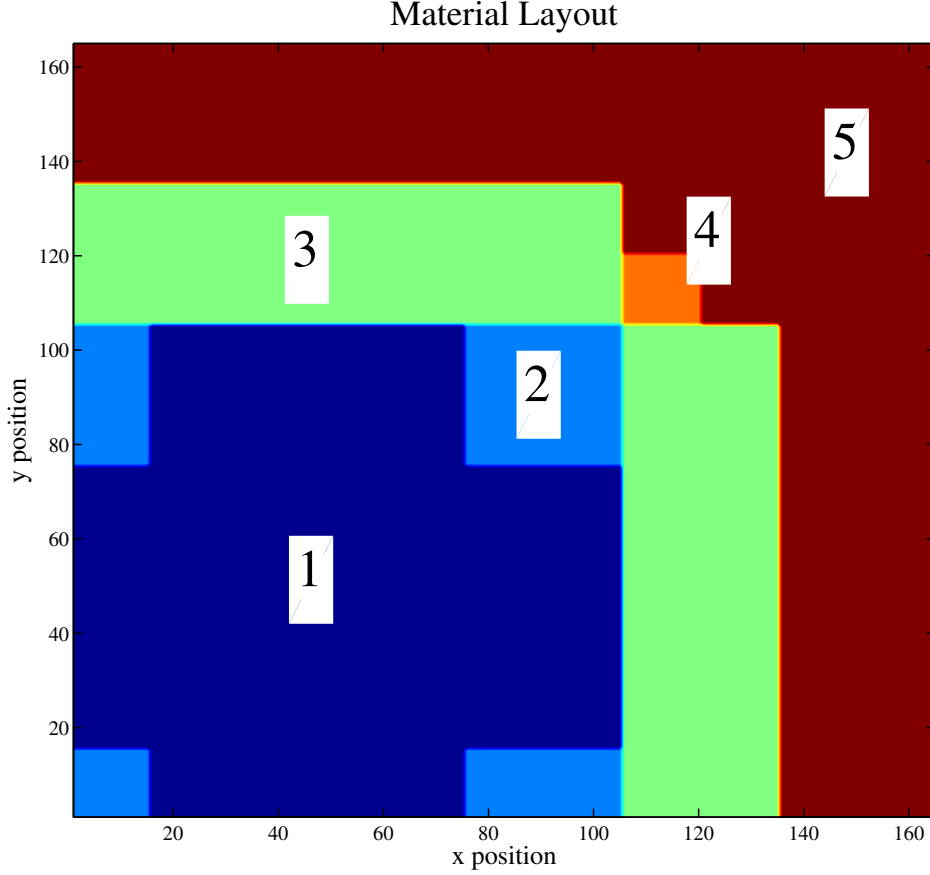


Figure 2: The material layout for the LRA-BWR 2-group eigenvalue problem.

low-order problem provides a scalar flux for transport sweep in the next iteration and a flat source is assumed in each cell. By assigning some shape to the source term, we hypothesize that we can at least partially remedy this error. Furthermore, in one spatial dimension, Willert showed that refining the spatial mesh allowed the user to quickly make the low-order discretization error smaller than the noise in the Monte Carlo simulation [3]. Not only can we notice a difference in the eigenvalue, but there is a noticeable difference between the deterministically computed eigenvector and the hybrid eigenvector, as displayed in Figure 3.

Table 1 lists the computed eigenvalue, hybrid and deterministic, and the

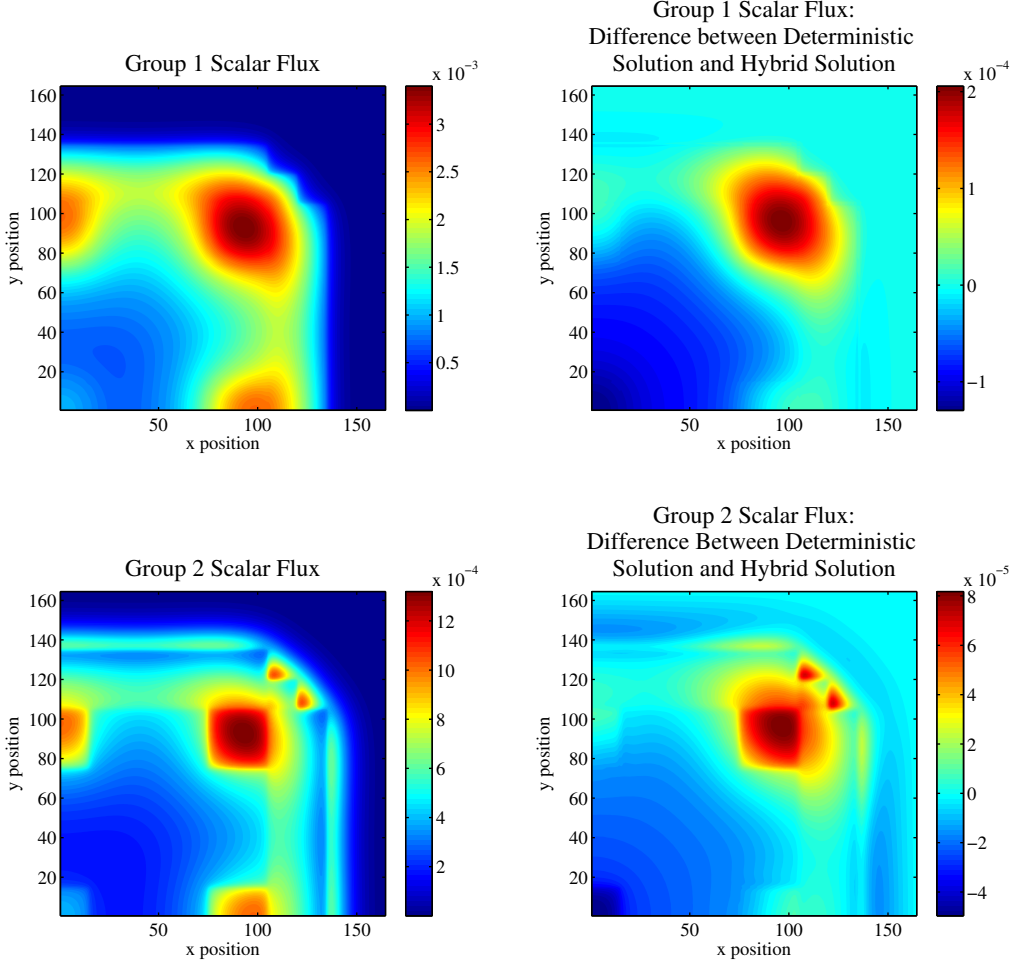


Figure 3: On the left we plot the eigenvector computed using the hybrid method. On the right we plot the relative difference between the hybrid and deterministic eigenvectors.

error compared to the analog Monte Carlo reference solution for a series of low-order spatial meshes. As we can clearly see, the error decreases as the mesh is refined. It is also important to note that the same number of particles are used in each simulation, regardless of the spatial mesh size.

We do notice, however, that the fission source distribution that is computed by the analog Monte Carlo simulation is far noisier than the fission

Table 1: Mesh Dependence of k -Eigenvalue for NDA-NCA(MC)

Mesh	Cell Width	Hyb. k_{eff}	Det. k_{eff}	Hyb. Error	Det. Error
88×88	$h = 1.875$	0.97816	0.97739	2.328E-02	2.405E-02
165×165	$h = 1.000$	1.00022	0.99896	1.220E-03	2.480E-03
220×220	$h = 0.750$	1.00059	0.99973	8.500E-04	1.710E-03
330×330	$h = 0.500$	1.00076	1.00041	6.800E-04	1.030E-03

source distribution computed via NDA-NCA(MC). This makes sense - analog Monte Carlo utilizes a very discrete method of tallying the fission source. Particles only contribute to the fission source tally if a fission interaction takes place in a given cell. On the other hand, within the NDA-NCA(MC) transport sweep, particles contribute to the scalar flux in each cell that they pass through. The comparison between these two quantities can be seen in Figure 4.

The convergence of the eigenvalue, eigenvector, low-order nonlinear residual and the Shannon entropy [18] can be visualized in Figure 5. After 15 iterations each sweep requires a total of roughly 2.5×10^{12} particle flights.

We can also analyze the convergence of the analog Monte Carlo calculation of the eigenvalue in Figure 6. As we can see, it requires at least 2000 active cycles to converge the eigenvalue to a relative tolerance of 10^{-5} . Each cycle utilizes 2×10^7 particles, each of which undergoes roughly 25 particle flights on average.

The NDA-NCA(MC) algorithm converges on the first five digits of the eigenvalue (compared to reference NDA-NCA(MC) solution) after roughly 6.4×10^{10} particle flights. The analog Monte Carlo algorithm required about 1×10^{12} particle flights to achieve the same degree of convergence. It is important to note that with the hybrid algorithm we have a much better mechanism for terminating the simulation. With the NDA-NCA(MC) algorithm we can use the low-order nonlinear residual, the change in the eigenvalue, the change in the eigenvector or the Shannon entropy as a convergence metric. While all of these quantities (aside from the low-order nonlinear residual) can be computed during the analog Monte Carlo simulation, we are more likely to encounter a false convergence event. At later stages in the analog Monte Carlo simulation, all of these values will be slowly changing, as they are computed via a running-average over active cycles.

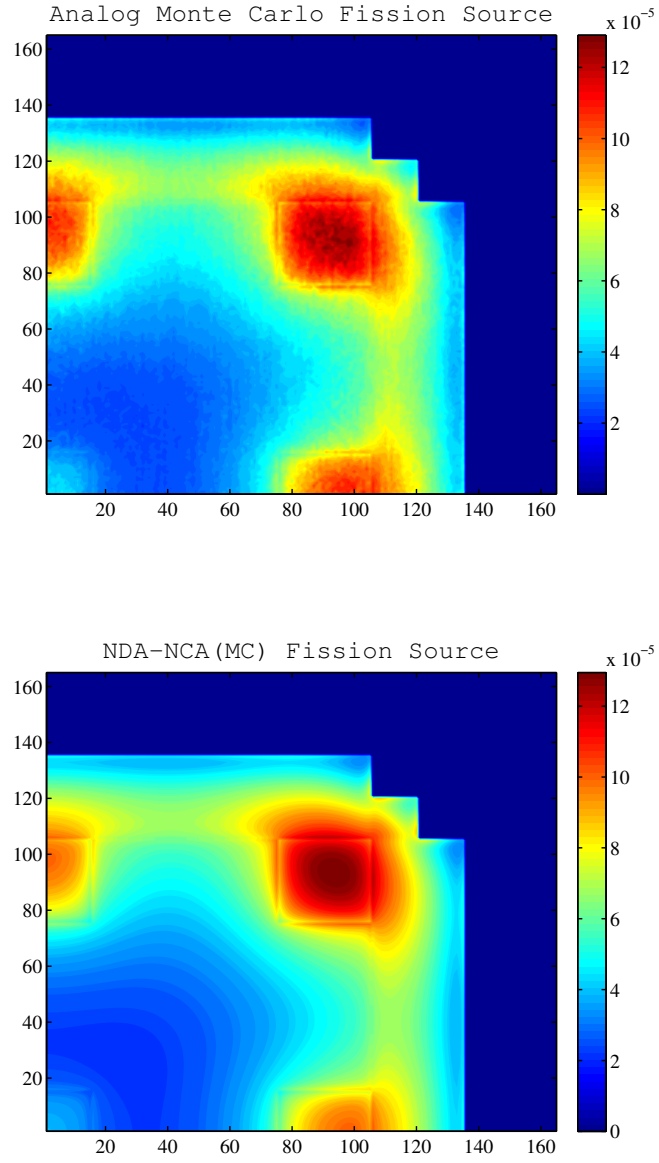


Figure 4: On the top we plot the fission source distribution computed using the analog Monte Carlo simulation. On the bottom we see the much less noisy fission source distribution computed using NDA-NCA(MC).

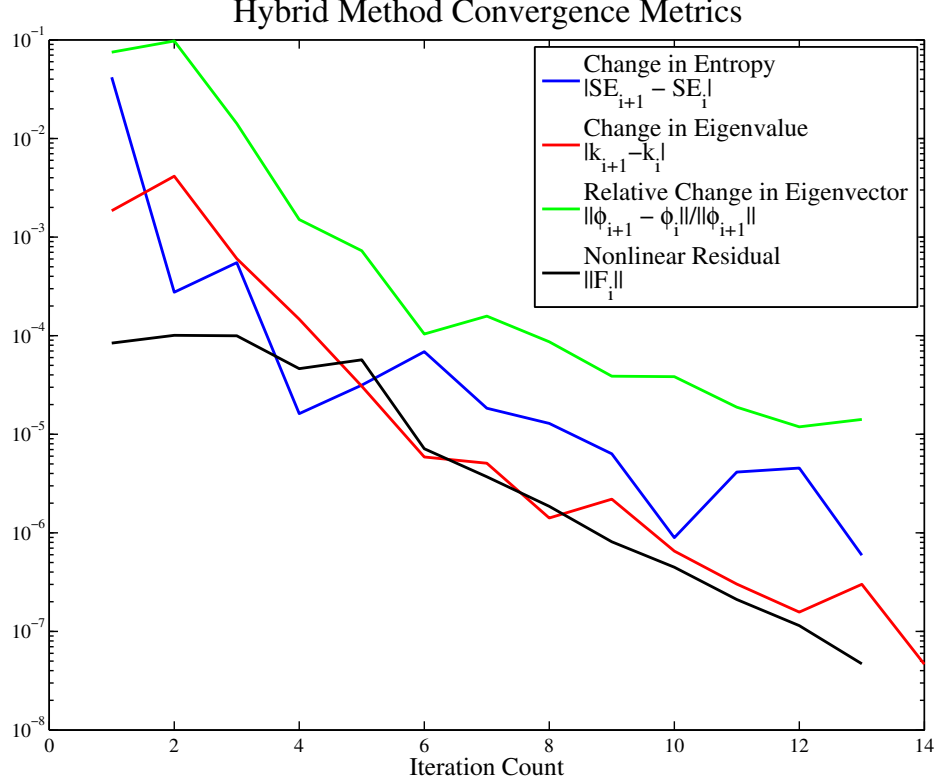


Figure 5: Convergence plots for NDA-NCA(MC) for the LRA-BWR problem.

6. Algorithmic Considerations

It is important to include a brief discussion on some of the more subtle differences between the analog Monte Carlo algorithm and the hybrid NDA-NCA(MC) algorithm. First of all, within the analog Monte Carlo algorithm all neutrons are born in the fissionable region. This is not the case when using the NDA-NCA(MC) algorithm; in this case, Monte Carlo particles are born everywhere throughout the domain to ensure adequate statistics for the scalar flux and current, which must be computed throughout the domain. This constitutes a major difference for a problem like the LRA-BWR test configuration. For this problem, roughly 65% of the cells contain fissionable material. This implies that the analog Monte Carlo algorithm gets to focus its generation of particles in a much smaller fraction of the domain. In

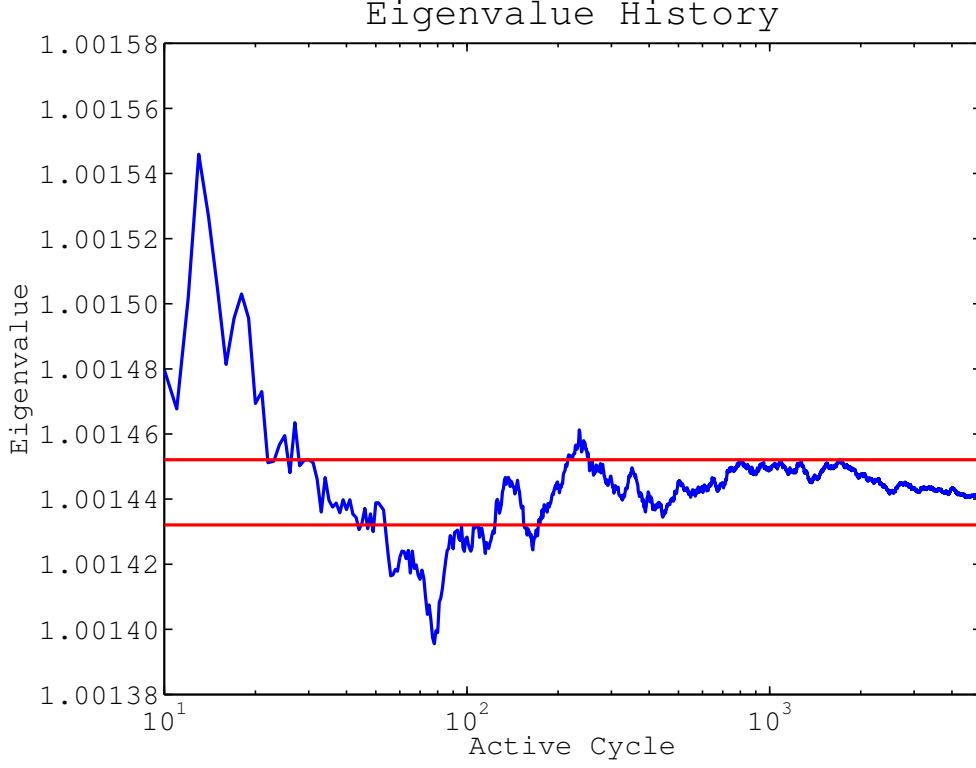


Figure 6: Convergence plot of the analog Monte Carlo simulation.

some sense, NDA-NCA(MC) is less efficient as particles must be distributed everywhere. On the other hand, analog Monte Carlo particles are inefficient as they only contribute to the fission source tally if they induce a fission event. NDA-NCA(MC) particles contribute to every cell and cell face which they cross as they stream from their birthplace to absorption location.

In the initial implementation of the NDA-NCA(MC) algorithm we chose a uniform distribution of particles throughout the spatial domain. We have proposed two methods for improving the algorithm by concentrating more particles in the fissionable area of the domain. We have implemented a non-uniform source which bases the number of particles in each cell on the relative source strength of that cell. This is a simple adaptation and does not introduce any additional sources of error, aside for the potential to achieve poor statistics in certain regions. Unfortunately, the difference in computational effort between the uniform and non-uniform particle source locations does

not appear to be significant. The second of these methods would involve a geometric hybrid in which the reflector region scalar flux and current would be computed using a deterministic algorithm. We leave this as an area of future work. This method has many appealing features, however may introduce new sources of difficult to quantify error.

7. Conclusion

We have developed a hybrid deterministic/Monte Carlo algorithm for solving the multi-group k -eigenvalue problem in multiple space dimensions. We have demonstrated the convergence of NDA-NCA(MC) for a 2-D, 2-group test problem and demonstrated that the hybrid method achieves more accurate results than the purely deterministic algorithm, however there is still some degree of low-order discretization error which corrupts the converged eigenvalue. Future work should consider higher-order discretization schemes for solving the consistent low-order diffusion eigenvalue problem in order to decrease the amount of error contributed from the low-order problem and should provide non-uniform cell-sources to the high-order problem.

Furthermore, we have demonstrated that the hybrid algorithm achieves a similar convergence rate of the eigenvalue in terms of total particle flights. However, we have also seen that the hybrid algorithm produces a far less noisy eigenvector and fission source distribution. If the fission source distribution is of high interest, using the hybrid algorithm may be a superior choice. Future work will involve adapting the NDA-NCA(MC) algorithm to utilize particle histories more efficiently.

References

- [1] C.T. Kelley, *Iterative Methods for Linear and Nonlinear Equations*, Society of Industrial and Applied Mathematics, Philadelphia, 1995.
- [2] D.A. Knoll, Kord Smith, and H. Park. *Application of the Jacobian-Free Newton-Krylov method to nonlinear acceleration of transport source iteration in slab geometry*, Nuclear Science and Engineering, 167(2):122-132, February 2011.
- [3] Jeffrey Willert, *Hybrid Deterministic/Monte Carlo Methods for Solving the Neutron Transport Equation and k -Eigenvalue Problem*, PhD Thesis, North Carolina State University, Raleigh, North Carolina, 2013.

- [4] Jeffrey Willert, C.T. Kelley, D.A. Knoll, and H. Park, *Hybrid Deterministic/Monte Carlo Neutronics*, SIAM J. Sci. Comput., 35(5), 62–83, 2013.
- [5] H. Park, D.A. Knoll, and C.K. Newman, *Nonlinear Acceleration of Transport Criticality Problems*, Nuclear Science and Engineering, 171:1-14, 2012.
- [6] D.A. Knoll, H. Park, and C. Newman, *Acceleration of k -Eigenvalue/Criticality Calculations using the Jacobian-Free Newton-Krylov Method*, Nuclear Science and Engineering, 167:133-140, 2011.
- [7] Min-Jae Lee, Han Gyu Joo, Deokung Lee and Kord Smith, *Investigation of CMFD Accelerated Monte Carlo Eigenvalue Calculation with Simplified Low Dimensional Multigroup Formulation*, PHYSOR 2010 - Advances in Reactor Physics to Power the Nuclear Renaissance.
- [8] Jinan Yang and Edward W. Larsen, *Calculations of k -Eigenvalues and Multi-Group Eigenfunctions Using the Hybrid “Functional Monte Carlo” Method*, PHYSOR 2010 - Advances in Reactor Physics to Power the Nuclear Renaissance.
- [9] Edward W. Larsen and Jinan Yang, *A Functional Monte Carlo Method for k -Eigenvalue Problems*, Nuclear Science and Engineering, 159:107-126, 2008.
- [10] Jeffrey Willert, C.T. Kelley, and D.A. Knoll, *Scalable Hybrid Deterministic/Monte Carlo Neutronics Simulations in Two Space Dimensions*, 2013 International Symposium on Distributed Computing and Applications to Business, Engineering, and Science, pp. 7-10, 2013.
- [11] Jeffrey Willert, C.T. Kelley, D.A. Knoll, and H.K. Park, *A Hybrid Approach to the Neutron Transport k -Eigenvalue Problem using NDA-based Algorithms*, Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering, 2013.
- [12] Jeffrey Willert and C.T. Kelley, *Efficient Solutions to the NDA-NCA Low-Order Eigenvalue Problem*, Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering, 2013.

- [13] Jeffrey Willert, X. Chen, and C.T. Kelley, *Newton's Method for Monte Carlo-based Residuals*, Submitted for publication, 2013.
- [14] K. S. Smith, *An Analytic Nodal Method for Solving the Two-group, Multidimensional Static and Transient Neutron Diffusion Equations*, Masters Thesis, Massachusetts Institute of Technology, 1979.
- [15] *Argonne Code Center: Benchmark problem book*, Argonne National Laboratory, Tech. Rep. ANL-7416, 1977, prepared by the Computational Benchmark Problems Committee of the Mathematics and Computational Division of the American Nuclear Society, Supplement 2.
- [16] J. Fleck and J. Cummings, *Implicit Monte Carlo Scheme for Calculating Time and Frequency Dependent Non-Linear Radiation Transport*, Journal of Computational Physics, 8:313-342, 1971.
- [17] E. E. Lewis and W. F. Miller, *Computational Methods of Neutron Transport*, American Nuclear Society, Inc., La Grange Park, 1993.
- [18] F. B. Brown, *On the Use of Shannon Entropy of the Fission Distribution for Assessing Convergence of Monte Carlo Criticality Calculations*, PHYSOR 2006 - ANS Topical Meeting on Reactor Physics Organized and Hosted by the Canadian Nuclear Society, Vancouver, BC, Canada, 2006.
- [19] Jeffrey Willert, C.T. Kelley, D. A. Knoll, D. Dong, M. Ravishankar, P. Sathre, M. Sullivan, W. Taitano, *Hybrid Deterministic/Monte Carlo Neutronics using GPU Accelerators*, in 2012 International Symposium on Distributed Computing and Applications to Business, Engineering, and Science, Q. Guo and C. Douglas, Eds. Los Alamitos, CA: IEEE, 2012, pp. 43-47.